

Please cancel claims 19-28 without prejudice.

- 1 1. (Original) A method for a storage operating system implemented in a storage system  
2 to concurrently perform readahead operations for a plurality of different read streams es-  
3 tablished in one or more files, directories, vdisks or luns stored in the storage system, the  
4 method comprising:  
5 receiving a client read request at the storage system, the client read request indi-  
6 cating client-requested data for the storage operating system to retrieve from a file, direc-  
7 tory, vdisk or lun stored in the storage system;  
8 determining whether the received client read request matches any of a plurality of  
9 readset data structures ("readsets") allocated for the file, directory, vdisk or lun contain-  
10 ing the client-requested data; and  
11 performing readahead operations in accordance with a set of readahead metadata  
12 stored in a readset that is determined to match the received client read request.
- 1 2. (Original) The method of claim 1, further comprising:  
2 allocating at least one readset for each of the one or more files, directories, vdisks  
3 or luns in which the plurality of different read streams is established;  
4 generating a separate set of readahead metadata for each of the plurality of differ-  
5 ent read streams; and  
6 storing each generated set of readahead metadata in a different readset allocated  
7 for the file, directory, vdisk or lun in which the read stream associated with the generated  
8 set of readahead metadata is established.
- 1 3. (Original) The method of claim 1, further comprising:  
2 initializing each allocated readset to store a predetermined set of values.

- 1 4. (Original) The method of claim 2, wherein the number of readsets allocated for a file,  
2 directory, vdisk or lun depends on the size of that file, directory, vdisk or lun.
- 1 5. (Original) The method of claim 4, wherein the number of readsets allocated for a file,  
2 directory, vdisk or lun is dynamically increased as the size of that file, directory, vdisk or  
3 lun is increased.
- 1 6. (Original) The method of claim 1, wherein a first readset is determined to match the  
2 received client read request if the first readset stores a set of readahead metadata associ-  
3 ated with a read stream that is extended by the client-requested data.
- 1 7. (Original) The method of claim 1, wherein a second readset is determined to match  
2 the received client read request when the client-requested data is located within a prede-  
3 termined fuzzy range associated with the second readset.
- 1 8. (Original) The method of claim 7, wherein the fuzzy range is derived based on a  
2 multiple of a number of client-requested data blocks specified in the received client read  
3 request.
- 1 9. (Original) The method of claim 7, wherein the fuzzy range extends in both a forward  
2 direction and a backward direction in relation to a last data block retrieved in a read  
3 stream associated with the second readset.
- 1 10. (Original) The method of claim 1, wherein a third readset is determined to  
2 match the received client read request if the third readset is determined to be unused.
- 1 11. (Original) The method of claim 10, wherein the third readset is determined to be un-  
2 used when a level value stored in the third readset equals a special indicator value.

1 12. (Original) The method of claim 1, wherein readahead operations are not performed if  
2 the storage operating system determines that the file, directory, vdisk or lun containing  
3 the client-requested data is accessed using a random access style.

1 13. (Original) The method of claim 12, wherein a DAFS cache hint included in  
2 the received client read request indicates that the file, directory, vdisk or lun containing  
3 the client-requested data is accessed using a random access style.

1 14. (Original) The method of claim 1, wherein readahead operations are not per-  
2 formed unless:

- 3 (i) a readset is determined to match the received client read request; and  
4 (ii) the matching readset stores a set of readahead metadata associated  
5 with a read stream that is extended by the client-requested data past a predeter-  
6 mined data block or memory address.

1 15. (Original) The method of claim 1, further comprising:  
2 if the received client read request does not match any of the readsets allocated for  
3 the file, directory, vdisk or lun containing the client-requested data, then performing the  
4 steps:

5 identifying the received client read request as being the first read  
6 request in a new read stream;

7 generating a set of readahead metadata associated with the new  
8 read stream;

9 selecting for reuse one of the readsets allocated for the file, direc-  
10 tory, vdisk or lun containing the client-requested data; and

11 storing the generated set of readahead metadata associated with the  
12 new read stream in the readset selected for reuse.

1 16. (Original) The method of claim 15, wherein the readset selected for reuse stores a  
2 level value that is less than or equal to level values stored in each of the other readsets  
3 associated with the file, directory, vdisk or lun containing the client-requested data.

1 17. (Original) The method of claim 1, wherein the client read request received at the  
2 storage system is a file-based client read request.

1 18. (Original) The method of claim 1, wherein the client read request received at  
2 the storage system is a block-based client read request.

19-28 (Cancelled)

1 29. (Original) A storage system that employs a storage operating system to concurrently  
2 perform readahead operations for a plurality of different read streams established in one  
3 or more files, directories, vdisks or luns stored in the storage system, the method com-  
4 prising:

5 means for receiving a client read request at the storage system, the client read re-  
6 quest indicating client-requested data for the storage operating system to retrieve from a  
7 file, directory, vdisk or lun stored in the storage system;

8 means for determining whether the received client read request matches any of a  
9 plurality of readset data structures ("readsets") allocated for the file, directory, vdisk or  
10 lun containing the client-requested data; and

11 means for performing readahead operations in accordance with a set of readahead  
12 metadata stored in a readset that is determined to match the received client read request.

1 30. (Original) A computer-readable media comprising instructions for execution in a  
2 processor for the practice of a method for a storage operating system implemented in a  
3 storage system to concurrently perform readahead operations for a plurality of different  
4 read streams established in one or more files, directories, vdisks or luns stored in the stor-  
5 age system, the method comprising:

6           receiving a client read request at the storage system, the client read request indi-  
7   cating client-requested data for the storage operating system to retrieve from a file, direc-  
8   tory, vdisk or lun stored in the storage system;  
9           determining whether the received client read request matches any of a plurality of  
10   readset data structures (“readsets”) allocated for the file, directory, vdisk or lun contain-  
11   ing the client-requested data; and  
12           performing readahead operations in accordance with a set of readahead metadata  
13   stored in a readset that is determined to match the received client read request.